

Analysis of Apple's Time Machine Backup Solution

Dain K. Schroeder
Office of Science, Community College Initiative Program

Big Bend Community College, Washington

Lawrence Berkeley National Laboratory
Berkeley, California

August 21, 2009

Prepared in partial fulfillment of the requirement of the Office of Science,
Department of Energy's Community College Initiative internship under the direction
of Charles E. Verboom in the Information Technology Division at Lawrence Berkeley
National Laboratory.

Participant: _____
Signature

Research Advisor: _____
Signature

ABSTRACT

Analysis of Apple's Time Machine Backup Solution. DAIN K SCHROEDER (Big Bend Community College, Moses Lake, WA 98837) CHARLIE E VERBOOM (Lawrence Berkeley National Laboratory, Berkeley, CA 94720).

At Lawrence Berkeley National Laboratory, it is common for a user to have unique or irreplaceable data on their work systems. As a result, backup and disaster recovery solutions are very important. According to a study published by Google, drives older than 2 years have up to a 7% failure rate. There are 2 options to deal with this problem: frequent backups before any loss of data occurs or expensive data recovery solutions when the drive has already failed. For institutional backups, Veritas NetBackup can be used. However, this solution has only recently been made compatible with Mac OS X Leopard, Apple's latest operating system. Apple's own backup utility called Time Machine is bundled with the operating system. Since Time Machine's introduction in 2007, LBNL has not studied its usefulness as an enterprise or workgroup solution. The information technology division chartered a project to recommend a set of services for users of Apple computers. This project is part of that effort. Tests were performed to compare storage devices (including USB, FireWire, Time Capsule, and an Apple XServe), as well as to compare Mac OS 10.5 with Mac OS 10.6. The study showed that although the results varied, Mac OS 10.6 is significantly faster at performing Time Machine backups to the same media. The study also showed that, in addition to operating systems, the selected target device also heavily impacts the time it takes to complete a Time Machine backup.

1. Introduction

Disaster recovery and backup solutions have been an important tool for computer users, and perhaps especially so in a Lab environment. It is not uncommon for an employee at Lawrence Berkeley National Laboratory (LBNL) to have irreplaceable or one-of-a-kind data stored on their work systems. According to a study published by Google¹, drives older than 2 years have up to a 7% failure rate each year. When a drive does fail, LBNL's Information Technology (IT) division recommends Kroll OnTrack, a professional data recovery company, for drive recovery. According to the OnTrack website, data loss can occur for several reasons, all of which can happen at the lab. These include human error, computer viruses, natural disasters, hardware and system problems, and software corruption or program issues.

There are two choices for dealing with a drive failure. The user can either restore from a backup (if one exists), or attempt an expensive drive recovery using a service like Kroll OnTrack. The IT division's solution, Veritas NetBackup, has supported PC operating systems, but has only recently been updated to support newer versions of the Mac operating system. Using Veritas, the IT division backs up around 70 terabytes (TB) of data (about 304 million files) each month from 422 hosts. The service may be more expensive than the user can afford and, for Mac users, it may not be as simple to use as native Mac software. In 2007 Apple provided a surprisingly elegant solution to this problem, known as Time Machine, for free with their latest operating system (Mac OS 10.5 "Leopard").

The IT division chartered a project to explore, assess, and recommend a set of IT services for users of Apple computers and iPhones at LBNL. Some of these services (such as file storage) will be visible to end-users. This included areas like data backup, which was the subject of this research project.

2. Methods and Materials

1. Theoretical Model

Time Machine is characterized as an incremental backup application. Typically, an incremental backup means that the software only backs up what's been changed since the last backup. Time Machine, however, differs from other incremental backups in one important way: before it begins to take incremental backups, it first does a backup of every file on the system.

When backing up over a network, this backup is made by creating a special type of intelligent disk image, called a sparse bundle. The unique thing about a sparse bundle when compared with a typical disk image is that the .sparsebundle file type has the ability to expand intelligently. In other words, while a typical 50 GB disk image (created with a product like Symantec Ghost) containing only 5 GB of real data would still take up 50 GB on the target hard drive, a sparse bundle only takes up the amount of space that it truly needs, meaning that it would take up 5 GB but have the ability to expand to 50 GB if more data was added. A .sparsebundle file is also somewhat unique in that it is not actually a single file; it is actually a package made up of many 8 megabyte (MB) pieces called bands. The benefit to this unique approach can be seen when a file in the image is removed or modified. Instead of having to re-copy everything that has already been backed up, Time Machine can

simply backup the relevant 8 MB pieces, with the result being that the backups complete faster. When the actual backup process occurs, the sparse bundle is mounted as a virtual drive on the client system.

When backing up to a local drive (via USB, FireWire, eSATA, or similar), Time Machine does not create a sparse bundle. Instead, it copies files from the local system directly into a folder on the destination drive called backups.backupdb. Since Time Machine is still not creating a bit-stream image (in other words, a complete hard drive copy), the backup will still only take as much space as is absolutely necessary. Because the backups are placed only in the backups.backupdb folder, the Time Machine backup drive may still be used to store data other than the backups.

Apple claims that in Mac OS 10.6 “Snow Leopard,” the next major release of their operating system, backups through Time Machine will be up to 50% faster.²

2. Apparatus

While many people who use Macintosh systems in their home networks love Time Machine for its ease of use, adequate testing has not been done to see how the relatively new backup solution performs in the enterprise. With an increasing number of scientists at LBNL using Macs, it has become important to recommend a course of action for backing up those systems. Over the course of the summer, extensive research was done into Time Machine. A lab environment was constructed to facilitate backing up to USB, FireWire, Apple Time Capsule, and a networked Time Machine server with a number of different configurations. These tests were done not only using Mac OS 10.5 Leopard (the most recent public release of Apple’s operating system), but also using the latest available beta seed of Mac OS 10.6 Snow

Leopard, which is set for a public final release sometime in September of 2009.³ In these tests, build 10A421a of the beta operating system was used.

3. Target Devices

Several devices were used to perform these tests. A 150GB CoolGear drive was used to perform tests over FireWire 400, FireWire 800, and USB 2.0, with additional USB 2.0 testing done on an Addonics SATA drive enclosure. Network backups to the Time Machine server were sent to an Apple Xserve with a 1TB SATA2 drive. See table 1 for a description of the backup media. See table 2 for a description of the systems used to perform these tests.

4. Experiment Setup

For these tests, two Mac Mini systems were used. One older system, the Mac Mini 1.1 model, was used for tests with USB 1.1 and FireWire 400. A newer Mac Mini was used for tests with USB 2.0 and FireWire 800. These systems were both prepared with what is considered to be a “standard Mac install,” including the default OS installation, iLife, FireFox, OpenOffice, and other common Mac applications. This “standard install” weighed in at very close to 30 GB.

Both Mac Mini systems were prepared using the standard procedure detailed above. After the machines were prepared, a custom script was used to launch a Time Machine backup to the media mentioned above while recording how long it took to complete that backup. Systems analysis tools were also employed to measure and report the impact the backup had on the relevant systems. Measurements were then taken on how long it took to backup to each media. Once all these tests were completed, the systems were updated to the latest available beta of Mac OS 10.6 in

order to conduct the same tests and determine the impact of the new operating system. Each of these tests was repeated multiple times on each operating system in an attempt to achieve consistent results.

To determine the speed difference between using Time Machine and manually copying files, a 30 GB disk image was also transferred to the same target devices. A custom-made Automator workflow fused with AppleScript was used to facilitate the automation of copies while timing the duration of the experiments. Network file transfers, however, were not supported via this workflow and were manually timed with a stopwatch iPhone app. These data transfer tests were also performed on both Mac OS 10.5 and 10.6.

To ensure that the backup tests were successful, selected backup tests were followed by a restore test to simulate a disaster recovery scenario. In an attempt to ensure a common starting point, the systems were formatted with Disk Utility prior to the restore. In each of these cases, the system could be successfully restored from the Time Machine backup.

3. Results

To understand the results of the tests, it is important to bear in mind what the maximum speed of each media is. USB 1.1 has a maximum transfer rate of 12 Mb/s, where USB 2.0 was significantly upgraded to 400 Mb/s. FireWire 400 can reach speeds of 400 Mb/s, while FireWire 800 can reach 800 Mb/s. All network tests were done over gigabit Ethernet, which has a maximum transfer rate of 1000 Mb/s. On Mac OS 10.5, FireWire 400 and FireWire 800 were the fastest device types tested, taking slightly over 3 hours to complete an initial backup. When the system

was updated to 10.6, FireWire 800 backups became significantly faster, and took almost an hour less to complete. It is possible that this occurred because Mac OS 10.5 is unable to utilize the full bandwidth of FireWire 800 for Time Machine, where Mac OS 10.6 might be. This is an area for further research. Running Mac OS 10.6, a server backup was the fastest, taking around 2.5 hours. See table 3 for the results of the Mac OS 10.5 experiments and table 4 for the Mac OS 10.6 experiments.

To perform backups to the Apple Time Capsule, three different sets of tests were done. First, systems backed up to the Time Capsule with a NetGear GS105 switch facilitating network traffic, which would have kept the devices from being affected by the LBNL subnet's bandwidth usage. Second, the Time Capsule was backed up to using a direct CAT6 connection in order to simulate ideal conditions. However, it was later discovered that this method might be inefficient. When a DHCP server cannot be located, the Time Capsule and Mac Mini would have fallen back to using Bonjour, Apple's implementation of Zero-Configuration Networking (ZeroConf). It is not known if using Bonjour would have had an adverse impact on the results. To account for this, a third set of tests were also done with the Time Capsule and host system on a private wired network together with a DHCP server and switch to facilitate network traffic.

The raw data copy tests were generally consistent, regardless of the client OS. The results that did have some variance involved were the server experiments. When the client OS was 10.5, copying to a 10.5 server took about 15 minutes and copying to a 10.6 server took 21 minutes. However, when the client was upgraded to 10.6, copying to a 10.5 server took 20 minutes and copying to a 10.6 server took 11

minutes. The fastest copies recorded were using FireWire 800, which took about 16 minutes to complete. For a full breakdown of results, see table 5 for information about OS 10.5 and table 6 for details on OS 10.6.

4. Discussion and Conclusions

As a result of these tests, it was discovered that the target device makes a distinguishable difference, as does the installed version of Mac OS X. On Mac OS 10.5 (Leopard), the fastest device tested for initial Time Machine backups was the FireWire 400 CoolGear drive. When the systems were updated to Mac OS Snow Leopard, it was observed that network file transfers to the server were significantly faster, as were server backups. In addition, initial Time Machine Backups were significantly faster using the USB 2.0 CoolGear drive. As indicated by figure 1, file transfers to a Snow Leopard server are faster with a 10.6 client than with a 10.5 client. This holds true for Time Machine backups as well. As figure 2 indicates, a 10.6 client finishes server backups faster than a 10.5 client, regardless of the server OS.

Another conclusion is that it is difficult to accurately predict how long an initial Time Machine backup may take, particularly when dealing with OS 10.5. On many of the tested devices (especially the two tested USB drives and the Time Capsule), the results varied greatly. This led the research group to believe that there may be some unknown or undocumented background process running, adversely affecting the results of the tests. Although the systems were closely monitored, it is also possible that a Spotlight system index may have taken place at the same time as some of the tests that took an abnormal amount of time.

In the restore tests that were done, Time Machine functioned as expected, and the restore times were significantly faster than backup times on the same media. For example, a backup to a Time Capsule that took 12 hours and 43 minutes was restored to the system in 2 hours and 43 minutes. Due to time constraints, it was not possible to do an in-depth study of Time Machine restores. It was discovered that each system had several disk permission issues after the restores the research group attempted. Users of Time Machine may wish to plan on spending some time repairing disk permissions with Disk Utility after a full Time Machine restore of their system.

For end users, Time Machine is very easy to use. While the backups take place, the background process (called 'backupd') does not interfere with the user and is only visible in the form of a small icon in the status bar, showing an animation of a clock turning backwards. To retrieve individual files, the user is shown an attractive space-themed screen, where there are simple 'back and forward' arrows allowing them to move 'back in time' to retrieve the file they want. To perform a full system restore, the user simply boots from the Mac OS installation media, selects 'Restore From Backup...' and follows the onscreen prompts. It should be noted that it is possible to restore from a backup of a different client system with apparently no issues.

5. Future Work

LBNL's Mac and PC support group will continue their ongoing effort to create and deploy services for Mac users around the Lab. An institutional backup solution

for the on-campus Macs will be recommended, which may be Time Machine or another solution (such as Veritas NetBackup).

Additional Time Machine research may also be done. A study of what causes the sometimes significant variability of the results would be beneficial in trying to estimate how long a standard backup might take. In addition, it would be beneficial to do a thorough study of backing up to a Mac OS X server located on a different subnet than the client.

6. Acknowledgements

This work was supported by the Center for Science and Engineering Education at Lawrence Berkeley National Laboratory, the Department of Energy, and the Office of Science. I would like to thank my mentor, Charles Verboom, as well as Big Bend Community College. I would also like to thank Ron Ustach at Apple for his valuable input, as well as the members of the www.macscripter.net forum for their friendly and expedient assistance in AppleScript automation of Time Machine.

7. References

- [1] "Failure Trends in a Large Disk Drive Population" Google
http://docs.google.com/gview?a=v&q=cache:q7nqEvSwbZII:labs.google.com/papers/disk_failures.pdf
- [2] "Mac OS X Snow Leopard – Refining the user experience" Apple
<http://www.apple.com/macosx/refinements/>
- [3] "Mac OS X Snow Leopard – The world's most advanced OS" Apple
<http://www.apple.com/macosx/>

Figures and Tables

<i>Device</i>	<i>Media Type(s)</i>	<i>RPM</i>	<i>Size</i>
<i>CoolGear Drive</i>	FW400, FW800, USB2	7200	150 GB
<i>Addonics Drive</i>	USB2	7200	120 GB
<i>Xserve</i>	SATA (internal)	7200	1 TB

Table 1 – Hardware specifications of drives used

<i>Device</i>	<i>Processor</i>	<i>RAM</i>	<i>HDD</i>
<i>Mac Mini 1.1</i>	1.83 GHZ Core Duo	1 GB 667 MHZ DDR2	80 GB
<i>Mac Mini 3.1</i>	2.0 GHZ Core 2 Duo	1 GB 1066 MHZ DDR3	120 GB
<i>XServe</i>	2.8 GHZ Quad Core Xeon	2 GB 800 MHZ DDR2	80 GB, 1 TB

Table 2 – Hardware specifications of systems used

<i>Client</i>	<i>Media / Target</i>	<i>Test 1</i>		<i>Test 2</i>		<i>Test 3</i>	
		<i>Hrs</i>	<i>Mins</i>	<i>Hrs</i>	<i>Mins</i>	<i>Hrs</i>	<i>Mins</i>
<i>Mac Mini 1.1</i>	<i>FireWire 400 CoolGear Drive</i>	3	7	3	8	3	43
<i>Mac Mini 3.1</i>	<i>FireWire 800 CoolGear Drive</i>	3	29	3	53		
<i>Mac Mini 1.1</i>	<i>USB 2.0 CoolGear Drive</i>	10	49	11	47		
<i>Mac Mini 3.1</i>	<i>USB 2.0 CoolGear Drive</i>	3	48	3	47		
<i>Mac Mini 3.1</i>	<i>USB 2.0 Addonics SATA enclosure</i>	3	31	24+ Hours ¹		12	40
<i>Mac Mini 1.1</i>	<i>USB 2.0 Addonics SATA enclosure</i>	3	27	9	6		
<i>Mac Mini 1.1</i>	<i>Time Capsule/Direct Ethernet</i>	12	43				
<i>Mac Mini 1.1</i>	<i>Time Capsule/LBNL Subnet</i>	4	39	13	39		
<i>Mac Mini 3.1</i>	<i>Time Capsule/LBNL Subnet</i>	12	1				
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.5, SATA 7200RPM</i>	4	12	4	14		
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.6, SATA 7200RPM</i>	8	25	4	10	4	11

1.) The backup took more than 24 hours to finish, and was cancelled.

Table 3 – Time taken for an initial Time Machine backup (OS 10.5.7)

<i>Client</i>	<i>Media / Target</i>	<i>Test 1</i>		<i>Test 2</i>		<i>Test 3</i>	
		<i>Hrs</i>	<i>Mins</i>	<i>Hrs</i>	<i>Mins</i>	<i>Hrs</i>	<i>Mins</i>
<i>Mac Mini 1.1</i>	<i>FireWire 400 CoolGear Drive</i>	3	26	2	38		
<i>Mac Mini 3.1</i>	<i>FireWire 800 CoolGear Drive</i>	2	30	2	29		
<i>Mac Mini 1.1</i>	<i>USB 2.0 CoolGear Drive</i>	2	32	2	32	2	33
<i>Mac Mini 3.1</i>	<i>USB 2.0 CoolGear Drive</i>	6	46	2	45		
<i>Mac Mini 3.1</i>	<i>USB 2.0 Addonics SATA enclosure</i>	6	20	2	43		
<i>Mac Mini 1.1</i>	<i>USB 2.0 Addonics SATA enclosure</i>	6	12	2	31		
<i>Mac Mini 1.1</i>	<i>Time Capsule/CAT6 Cable</i>	3	27	3	25		
<i>Mac Mini 1.1</i>	<i>Time Capsule/LBNL Subnet</i>	6	36				
<i>Mac Mini 3.1</i>	<i>Time Capsule/LBNL Subnet</i>	6	54	3	26	3	31
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.5, SATA 7200RPM</i>	2	42	2	48		
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.6, SATA 7200RPM</i>	2	38	2	42		

Table 4 – Time taken for an initial Time Machine backup (OS 10.6 beta)

<i>Client</i>	<i>Media / Target</i>	<i>Test 1</i>	<i>Test 2</i>
		<i>Minutes</i>	<i>Minutes</i>
<i>Mac Mini 1.1</i>	<i>FireWire 400 CoolGear Drive</i>	31	31
<i>Mac Mini 3.1</i>	<i>FireWire 800 CoolGear Drive</i>	15	16
<i>Mac Mini 1.1</i>	<i>USB 2.0 CoolGear Drive</i>	33	28
<i>Mac Mini 3.1</i>	<i>USB 2.0 CoolGear Drive</i>	23	22
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.5, SATA 7200RPM</i>	14	15
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.6, SATA 7200RPM</i>	21	14

Table 5 – Time taken to manually copy a 30 GB image file (10.5 client)

<i>Client</i>	<i>Media / Target</i>	<i>Test 1</i>	<i>Test 2</i>
		<i>Minutes</i>	<i>Minutes</i>
<i>Mac Mini 1.1</i>	<i>FireWire 400 CoolGear Drive</i>	33	32
<i>Mac Mini 3.1</i>	<i>FireWire 800 CoolGear Drive</i>	16	16
<i>Mac Mini 1.1</i>	<i>USB 2.0 CoolGear Drive</i>	39	42
<i>Mac Mini 3.1</i>	<i>USB 2.0 CoolGear Drive</i>	23	25
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.5, SATA 7200RPM</i>	20	19
<i>Mac Mini 1.1</i>	<i>Xserve, OS 10.6, SATA 7200RPM</i>	11	12

Table 6 – Time taken to manually copy a 30 GB image file (10.6 client)




Figure 1 – Minutes taken to copy a 30 GB image file over AFP




Figure 2 – Hrs:Mins to complete an initial Time Machine backup